

Method and Apparatus for Determining Latency Between Multiple Servers and a Client

5

CROSS-REFERENCES TO RELATED APPLICATIONS

The present application claims priority to U.S. Provisional Application No. 60/193,988 filed March 31, 2000 (Attorney Docket No. 4878-US), commonly owned, and hereby incorporated by reference for all purposes.

BACKGROUND OF THE INVENTION

15 TECHNICAL FIELD

The invention relates to the routing of requests to a networked server in a computer environment. More particularly, the invention relates to determining a dynamic hop count between two nodes across a network in a computer environment.

DESCRIPTION OF THE PRIOR ART

The Internet is a confederation of loosely connected networks that connect computers all around the world. The proliferation of information and immense user population has made the management of information critical. Information is mirrored on multiple servers in order to improve performance, availability, and scalability.

To facilitate the mirroring of information, clients requesting information need to be routed to the optimal server. The traffic management industry routes traffic to the optimal server using different latency metrics including the Round Trip Time (RTT) and hop count. The Round Trip Time measures the time it takes for a
5 packet to travel between a server and a client or another server.

The dynamic hop count between a client and a server is typically derived from the Border Gateway Protocol (BGP). BGP is specified in RFC 1267 published in October 1991 and RFC 1654 published in July 1994. BGP gives the hops
10 between different Autonomous System Numbers (ASN). The dynamic hop count can be used to differentiate latency metrics that might be close in terms of RTT.

The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. The network reachability
15 information includes information on the Autonomous Systems (AS) that the reachability information traverses. A BGP speaker advertises to its peers, *i.e.*, other BGP speakers that it communicates with, in neighboring ASs only those routes that it uses. The information is sufficient to construct a graph of AS connectivity from which routing loops may be pruned.
20

Determining the dynamic hop count is currently problematic due to the following reasons:

1. **BGP hops:** The actual number of hops will typically be larger than that
25 advertised by the BGP protocol since BGP only gives the hops between the ASNs.
2. **IP hops:** Getting hop counts from the Internet Protocol (IP) is difficult because the Time To Live (TTL) field that is part of the IP header is not initialized to standard values by the various TCP/IP stack software running
30 on various Operating Systems, *e.g.*, Windows 98, Linux, NT, etc.

The invention described in this patent provides a different and more precise method of determining the dynamic hop count.

- 5 It would be advantageous to provide a method and apparatus for determining latency between multiple servers and a client that provides a more precise method of determining dynamic hop counts. It would further be advantageous to provide a method and apparatus for determining latency between multiple servers and a client that reduces the traffic required to measure the hops across
10 the network.

SUMMARY OF THE INVENTION

- 15 The invention provides a method and apparatus for determining latency between multiple servers and a client. The system provides a more precise method of determining dynamic hop counts and optimal content servers. In addition, the invention reduces network traffic required for measuring the dynamic hop counts.
20 A preferred embodiment of the invention receives requests for content server addresses from local domain names servers (LDNS). POPs that can serve the content are determined and sent latency metric requests.

25 The content server receives the request for latency metrics and looks up the latency metric for the client of the requesting LDNS. Periodic latency probes are sent to the IP addresses in a Latency Management Table. The IP addresses of clients are masked so the latency probes are sent to higher level servers to reduce traffic across the network. The hop count and latency data in the packets sent in response to the latency probes are stored in the Latency Management
30 Table.

The information in the Latency Management Table is used to determine the latency metric from the resident POP to the requesting client before sending the latency metric to the requesting server. The BGP hop count in the Latency Management Table is used for the latency metric upon the first request for an IP
5 address. The latency metric is calculated for subsequent requests of IP addresses using the hop count and RTT data in the Latency Management Table.

Latency metrics from POPs are collected and the inverse relationship of the hop counts in a weighted combination with the RTT are used to determine which
10 latency metric indicates the optimal POP. The address of the optimal POP is then sent to the requesting LDNS.

Other aspects and advantages of the invention will become apparent from the following detailed description in combination with the accompanying drawings,
15 illustrating, by way of example, the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

20 Fig. 1 is a block schematic diagram of a preferred embodiment of the invention measuring latency between servers and a client according to the invention;

Fig. 2 is a block schematic diagram of an example of a preferred embodiment of the invention measuring dynamic hop counts from two POPs to a Border
25 Gateway server according to the invention;

Fig. 3 is a diagram of a Latency Management table according to the invention;

Fig. 4 is a block schematic diagram of an example of differing TTL values in IP
30 packets according to the invention;

Fig. 5 is a block schematic diagram of the inverse relationship of dynamic hop counts in a preferred embodiment of the invention according to the invention; and

- 5 Fig. 6 is a block schematic diagram of a task-level viewpoint of a preferred embodiment of the invention according to the invention.

DETAILED DESCRIPTION OF THE INVENTION

10

The invention is embodied in a method and apparatus for determining latency between multiple servers and a client in a computer environment. A system according to the invention provides a more precise method of determining dynamic hop counts and optimal content servers. In addition, the invention 15 provides a system that reduces network traffic required for measuring the dynamic hop counts.

The invention provides a new method to determine the dynamic hop count between two nodes (client and server). This new method provides a dynamic 20 hop count that is more precise than the hop count obtained using the Border Gateway Protocol (BGP).

The latency probes in a Speedera Network are responsible for determining the latency between the Speedera servers and a client. This latency information is 25 used by the Speedera Domain Name Server (SPDNS) to direct a client to the server that is "closest" to the client in latency.

Referring to Fig. 1, each Speedera Point of Presence (POP) 101, 102 has a probe server 104, 106 and other infrastructure servers 103, 105. When the 30 Client 109 receives a request, it tries to resolve the request from the Local Domain Name Server (LDNS) 108. If the LDNS 108 cannot resolve the request,

it forwards the request to the Speedera Domain Name Server (SPDNS) 107. The SPDNS 107 then routes the request to the optimal POP 101, 102 by determining the server “closest” to the LDNS 108.

- 5 In this example, the Client 109 performs a name lookup for web content images.rich.com/images 110. The local DNS (LDNS) client 108 forwards the request 111 to the SPDNS 107.

The SPDNS 107 requests latency information 112, 113 from the Speedera
10 probes 104, 106 at locations that can serve images.rich.com (POP1 101 and
POP2 102). The latency probes 104, 106 initiate probes 117, 118 to the LDNS
108. The latency probes 104, 106 return the latency metrics 115, 116 to the
SPDNS 107.

- 15 The main components that are used by the invention to determine latency
metrics are as follows:

- Sub A 7
20 • RTT time from the latency probe to LDNS.
• ASN (Autonomous System Number) routing information derived from
Border Gateway Protocol (BGP).
• Dynamic hop count from the latency probe to LDNS.

BGP is a standard algorithm implemented in routers. It is a routing protocol that
is used between large routers covering large administrative domains. All routes
25 that are available within a network are exported to another network in an
abbreviated form. The network reachability information includes information on
the Autonomous Systems (AS) that the reachability information traverses. A
BGP speaker (router) advertises to its peers, i.e., other BGP speakers that it
communicates with, in neighboring ASs only those routes that it uses.

An AS is a set of routers under a single technical administration, using an interior gateway protocol and common metrics to route packets within the AS, and using an exterior gateway protocol to route packets to other ASs.

Sub A²

5 The latency probe uses a UDP Reverse Name Lookup and Traceroute to determine RTT and dynamic hop count. Reverse Name Lookup is a standard DNS query that specifies a client IP address and asks for the client name. Traceroute is a specific format of a packet that is sent between routers to indicate if a packet has reached a destination.

10

Obtaining absolute hop counts is ideal, however, this is not a realistic goal in the real world. Quite often some of the local DNSs are unavailable because they are sitting behind a firewall or the DNS will drop the packet thinking that it is a denial of service attack. With these restrictions, it is not possible to reliably use an 15 absolute hop count. Additionally, as described below, the IP packet information is not uniform throughout the network.

With respect to Fig. 2, most of the time the ASN 203 is easily reached. Even from various diverse points it is possible to converge onto the AS 203. However, 20 it is the path required to reach the AS 203 from each server that is the important factor.

POP1 201 and Pop2 202 must each find the distance from themselves to the ASN 203. The latency for each path must also be found. The distance and 25 latency are calculated by using the hop count and latency time, respectively.

Currently, it is only possible to measure hop count and latency time by sending a packet to the destination. However, as previously mentioned, most of the time the packet is sent back as not being able to reach the destination.

30

The invention aggregates the client 208 and the DNS 207 and assumes that they are co-located. Additionally, the latency and the hop counts are measured up to the Border Gateway (BG) 206. Once the autonomous system is entered, the hop counts are not as important. The distance from the BG 206 to the client 208
5 is the same from either POP 201, 202 at that point. Therefore, the relevant distance is to the BG 206. In other words, the distance T1 210 and T2 211 are most likely not equal, but the distance T3 212 is the same for both POPs 201,
202.

~~Sub A3~~ 10 The address of the client 208 is masked to the IP prefix of the ASN 203. For example, if the address of the client 208 is 4.10.20.30 then the address is masked to the ASN 203 which is 4.0.0.0. The mask can vary depending on the granularity desired. For example, the first eight bits may be masked to the DNS, i.e., the address would then be masked to 4.10.20.00. This can be adjusted
15 depending on the size of the network.

The data in the BGP table does not change over a period of time. Incremental updates are sent as the routing tables change. However, BGP does not require periodic refresh of the entire BGP routing table. Therefore, a BGP speaker must
20 retain the current version of the entire BGP routing tables of all of its peers for the duration of the connection.

The invention performs its measurements a few times a day to achieve a good picture of the network status. The invention reduces Internet traffic by
25 aggregating at a higher level and therefore requires less probes.

~~Sub A4~~ Referring to Fig. 3, a latency management table is used by the invention. The table contains the following fields: IP address 301; BGP hop count 302; and Trace data 303. The IP address field 301 contains the IP addresses that the
30 server is responsible for. BGP hop counts 302 is taken from the BGP routing

~~Sub A⁴~~ table for the particular IP address. The Trace data field 303 contains the latency and hop count information obtained by the latency probes.

~~Sub A⁵~~ The BGP hop count 302 is used by the SPDNS when the first request from an LDNS comes in for a particular IP address. No previous connection has been established at this point. This is because the dynamic hop count to the BG takes some time to actually measure. Subsequent requests use the actual BG hop count measured by the system located in the Trace field 303. The latency measurement is also taken and the combination of the hop count and latency measurement in the Trace field 303 is used to determine which SPDNS is closest to the client.

~~Sub A⁶~~ The invention does not need to get absolute hop counts from the various probe servers to the LDNS. It is sufficient to determine the relative hop counts between the probe servers and the LDNS and use this information in arriving at relative latency metrics between the client and the probe servers. This information can then be used by the Speedexa DNS to choose between the service points and locate a server that is closest to the client.

With respect to Fig. 4, the invention is based on the very safe assumption that the target LDNS 402 will always send out packets with a fixed TTL independent of the SPDNS 401 that initiated a probe. Getting hop counts from the Internet Protocol (IP) is difficult because the Time To Live (TTL) field that is part of the IP header is not initialized to standard values by the various TCP/IP stack software running on various Operating Systems, e.g., Windows 98, Linux, NT, etc.

In this example, the SPDNS 401 originates a packet with a TTL of 64 403. The LDNS 402 always sends out packets with its own independent TTL value of 128 404. DNS's will send whatever TTL value that they prefer. The values may differ between each DNS, but each DNS is consistent in its TTL value.

Sub A7

Referring to Fig. 5, the two probe servers (Probe1 502 and Probe2 503) have initiated probes to the LDNS 501. The LDNS 501 returns a response to each of the probe servers. Here, the LDNS 501 sets the TTL in the response IP packet to R 504, 505. The TTL will be decremented by 1 each time the packet passes through a router (hop). In this example, the packets go through $H1$ hops 506 between LDNS 501 and Probe1 502 and $H2$ hops 507 between LDNS 501 and Probe2 503. The TTL of the response packet that arrives at Probe1 502 will be $(R-H1)$ and which arrives at Probe2 503 will be $(R-H2)$. The TTL of the response at Probe1 502 and Probe2 503 will be in inverse relation to the number of hops 508, 509 (The fewer the number of hops, the higher the TTL). Thus, a relative hop count can be arrived at by subtracting the TTL in the response packet from a fixed value.

Sub A8

The latency metric is a weighted combination of the RTT and the hop count. The latency metric is used to determine the server that is the most efficient for accessing a client. The invention precisely determines the hop count metric between client and server.

The dynamic hop count metric derived through the invention is more accurate than the hop count derived from BGP. As a result, requests will be routed more accurately to the optimal server resulting in improved performance. The invention improves performance in multiple Internet infrastructure products including performance monitoring, WAN traffic management, and content distribution.

With respect to Fig. 6, a task level viewpoint of the invention is shown. Requests for content server addresses from LDNS's are received by the Receive IP Address Request module 605. The Receive IP Address Request module 605 sends the content request to the Request Latency Metrics module 606. POPs that can serve the content are retrieved from the Server Table 609 by the Request Latency Metrics module 606. The Request Latency Metrics module 606

then sends latency metric requests to the POPs that can serve the content and notifies the Determine Optimal Server module 608 that latency metrics are expected.

Sub A 5 Request for latency metrics activate the Send Latency Metric module 601 to lookup the latency metric for the requesting LDNS's client. The Send Latency Probe module 603 sends latency probes to the IP addresses in the Latency Management Table 604. The IP addresses of clients are masked so the latency probes are sent to higher level servers as detailed above. Packets sent in 10 response to the latency probes sent by the Send Latency Probe module 603 are received by the Receive Response Packet module 602. Hop count and latency data are stored into the Latency Management Table 604.

Sub A 15 The Send Latency Metric module 601 uses the information in the Latency Management Table 604 to determine the latency metric from the resident POP to the requesting LDNS's client before sending the latency metric to the requesting server. The Send Latency Metric module 601 uses the BGP hop count in the Latency Management Table 604 for its calculations upon the first request for an IP address. The latency metric is calculated for subsequent requests of IP 20 addresses by the Send Latency Metric module 601 using the hop count and RTT data obtained from the Receive Response Packet module 602.

Sub A 25 Latency metrics from POPs are received by the Receive Latency Metrics module 607. The latency metrics are sent to the Determine Optimal Server module 608. The Determine Optimal Server module 608 gathers the expected latency metrics and uses the inverse relationship of the hop counts in a weighted combination with the RTT to determine which latency metric indicates the optimal POP. The Determine Optimal Server module 608 then sends the address of the optimal POP to the requesting LDNS.

Although the invention is described herein with reference to the preferred embodiment, one skilled in the art will readily appreciate that other applications may be substituted for those set forth herein without departing from the spirit and scope of the present invention. Accordingly, the invention should only be limited by the Claims included below.